Enabling Harbor to Harbor Autonomous Situational Awareness in Sea Ice Conditions [ENHANCE]

Finnish Geospatial Research Institute Aalto University Fleetrange OY

June 16, 2023

Artificial intelligence (AI) and machine learning (ML) algorithms in collaboration with an assembly of environment perception sensors provide an autonomous vessel with precise location and situational awareness. However, studies on traditional sensor assembly and related AI techniques and their performance under wintertime sea and weather conditions is still scarce. Especially, the prevalent sea-ice conditions need to be detected via on-board Red Green Blue (RGB) and infra-red cameras under all-weather and visibility conditions as part of the autonomous situational awareness process. Secondly, the resilience of the vessel's systems requiring GNSS signals needs to be studies under winter conditions. For example the impact of multipath effects caused by sea-ice needs to be better understood.

Autonomous shipping is currently foreseen to include two primary types of operation, remote control or fully autonomous. However, a hybrid solution in which the ships can operate fully autonomously in open waters and taken under remote control for docking and harbour operations is likely the first step in autonomous operations. End-to-end fully autonomous shipping with auto-docking is also an option, but mostly this would be used for heavy frequency crossing ferries or "Yara Birkeland", the first fully electric and autonomous container vessel in the world, type of specialised point-to-point operations where the port infrastructure can be custom built to support an autonomous vessel.

Applications for special ice navigation sensing techniques can be utilised for example in the following tasks:

- In the first stage an intelligent and autonomous sensing can benefit human operators on conventional ships in the form of an increased situational awareness on-board the vessel
- Collecting this data can be fed to situational awareness systems ashore or to weather observation platforms to improve ice forecasting and ice chart creation.
- The data is valuable as machine learning datasets for further system development
- In the second stage the data can be directly used by the autonomous vessels control systems to improve the vessels capabilities when operating in ice conditions

Ice navigation and navigation where ice sensing is important is done primarily in four geographical areas: The Baltic Sea, The Great Lakes, The North Sea Route (NSR), and the polar regions (Arctic and Antarctic). Of these areas the Baltic Sea and the Great Lakes see heavy goods and passenger traffic, while the NSR only sees very limited goods traffic. The Arctic and Antarctic areas are mostly passenger vessel cruising grounds and operated mostly in the summer months.

Winter navigation places higher requirements on the vessel, such as a stronger hull, higher engine power and better manoeuvrability. Furthermore, ice going ships are usually also constructed with higher levels of redundancy. In addition to the technical aspects, ice going ships are also manned by a more skilled, and thus expensive, crews. This all translates to higher capital expenditure (CAPEX) and operating expenditure (OPEX), which means that the underlying business case must be more profitable than what is typical for open water operations in warm climates.

The design philosophy for ice going ships is much the same as for autonomous vessels, which also require higher levels of redundancy, better manoeuvrability, and better sensing equipment. We can therefore assume that constructing and operating ice going autonomous vessels may even be more viable than doing the same for conventional non-ice going ships, since the difference in CAPEX may be less than what would be the case for conventional non-ice going ships.

One of the critical challenges of winter navigation is the environment. Visibility is often restricted due to rain, snow or fog and the lighting conditions are difficult with either no light or a very low-lying sun that causes reflections and extreme variations on light and contrast. Outside air and water temperatures, together with humidity, rain, sleet, snow and waves, affect the ship through icing of equipment and superstructures and in worst cases can lead to capsizing of the vessel (due to high centre of gravity). In most cases it causes degradation in performance of sensing equipment. The winter environment also includes sea ice, which is difficult to predict, can shift and move extremely fast and which causes level ice fields, drift ice, pack ice and ridges. Ice fields also create pressure on the hull and can lead to significant hull and propeller damages. Thus careful planning is required when preparing the ship to operate in icy conditions. This includes better protection of the electronics and heating of externally placed sensors and systems and to some extent also specialised capabilities, such as the ability to see in the dark or read ice data with radars or from satellite images.

Operational factors include voyage planning, considering other traffic, understanding the capabilities and restrictions of manoeuvrability in ice and ice tracks, knowing how the icebreaker and traffic management works in winter conditions, along with understanding the concepts of icebreaker assistance.

Despite the extensive guidelines and regulations related to winter and cold weather maritime operations, there is very little publicly available data that would be suited to explore the feasibility of AI or machine learning concepts for autonomous navigation or increasing situational awareness. The current state-of-the-art AI algorithms have already been shown to enable autonomous driving of cars and it is likely that similar algorithms could be used in maritime environment. The ship use case should also be simpler compared to cars due to being much slower paced.

Serious attempts at commercialising autonomous driving were started years ago, and the difference in maturity compared to autonomous ships in the research infrastructure is evident. There are publicly available datasets for all different tasks related of autonomous driving, for example pre-trained neural network models. The existing data and tools allows researchers to experiment quickly and to build upon the previous state-of-the-art. A significant part of the publicly available autonomous driving data is in the form of RGB images, but there is also a respectable amount of data from LiDARs, IR cameras, and even radars. The amount of publicly available resources related autonomous ships is limited in comparison.

The ENHANCE (Enabling Harbor to Harbor Autonomous Situational Awareness in Sea Ice Conditions) project is aimed to study How to improve the autonomous navigation and situational awareness in the maritime domain by utilising ML and AI concepts. For this end, we utilise both traditional RGB cameras as well as IR cameras to obtain as diverse as possible data set of sea-ice conditions to train a ML model that can identify an ice tracks and objects, such as buoys, that can help in navigation. In addition to the camera concepts, we study the impact of possible multipath impacts caused by sea-ice on traditional GNSS receivers. As with the ML data sets, the number of studies focusing on behaviour of the GNSS signals in maritime winter navigation are limited, thus additional information is needed to better understand possible problems that will be encountered when designing a fully autonomous vessel.

1 System designs and data analysis methods

Because of the limited publicly available data for machine learning purposes, a data collection campaign was required. The cruise vessel MS Megastar, owned by the company Tallink, had been utilised as a test platform during previous maritime project (Maritime AI-NAV), and was naturally chosen as a suitable test platform also for the ENHANCE project. Unfortunately, the sea-ice conditions in the Gulf of Finland were not suitable to obtain sufficient amount of data. The winter of our field data campaign was mild and the ice coverage near Helsinki was limited. Thus, secondary test campaigns were organised, utilising the icebreaker Sampo, operating in the Bay of Bothnia and a ferry operating between Helsinki and the island of Suomenlinna.

The antenna used for the campaign was a Novatel-GPS-703-GGG-HV, with the capability of receiving at L1/E1, L2, L5 and E5a/b channels. The antenna was connected via a power divider/amplifier which provided power for the antenna and divided the incoming signal to three different devices. The main receiver was a Novatel ProPak6 which was used to measure constantly the trajectory of the vessel. The two other devices connected to the amplifier were a u-blox M8T receiver and a LabSat3 record and replay device. The M8T receiver was also connected to a laptop running the 'u-centre' software recording the receiver output. The LabSat3 was set to record the L1/E1 and E5a bands as a backup in case there would have been problems with the Novatel receiver so that the campaign could be re-played in laboratory.

The camera set-up on-board Sampo Icebreaker consisted of two FLIR cameras: FLIR M232 infrared maritime camera, and FLIR Backfly S visible range RGB camera. These cameras were placed rigidly to enable data fusion of the RGB and IR images, however, this meant that the cameras could not be moved after installation. The goal on-board the Sampo icebreaker was simply to collect as much data as possible to ensure that we would have enough samples for different types of ice. All of this data will then act as either training or testing data for the machine learning methods.

The first step in the fused RGB/IR prediction is to align the images on a pixel level. The alignment is done by estimating an affine or projective transform between a pair of RGB and IR images taken at the same time. Estimating this type of affine (or projective) transforms between images is standard in computer vision application given a matching set of points from the images, with well-known methods for extracting these matching set of points automatically. However, these methods are meant for the cases where the images are of the same spectra. These automatic point extraction and matching methods did not work for a RGB and IR image pair, thus the point selection and matching was done manually. The estimated

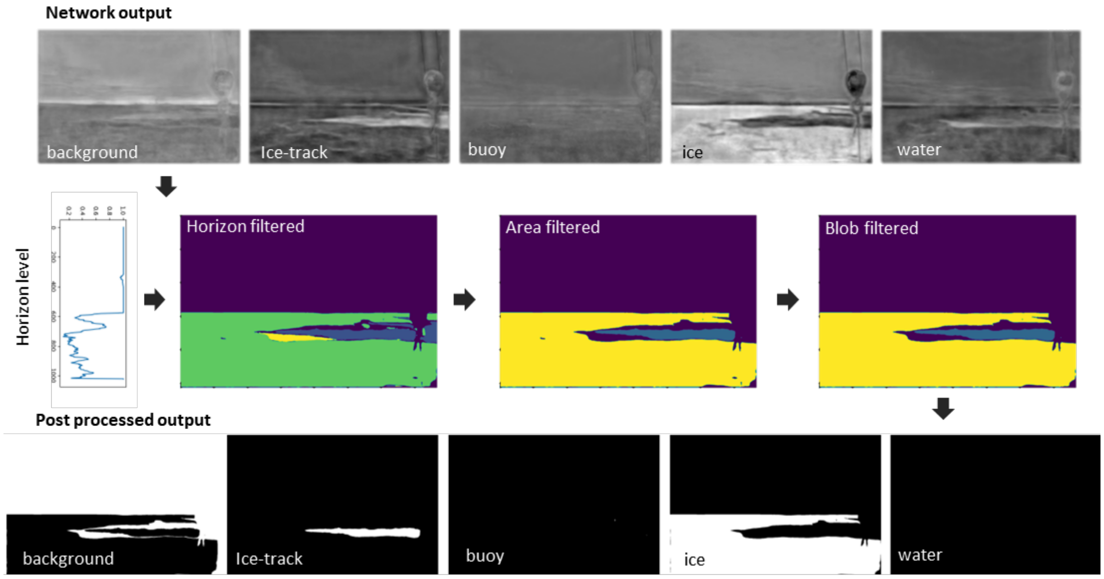

Figure 1: Examples of the semantic segmentation and the post-processing of a RGB image. The first row shows the raw network output, the second row shows the filters to be applied, and the third row shows the output of post processing.

transform is valid as long as the relative transformation between the RGB and IR camera does not change, but making this process automatic is something that could be explored in the future.

The resulting transform can be used to map the pixels in the RGB image to the IR image frame, thus the images are transformed to the same coordinate system. This allows us to translate detection from one image to the other, which is required for data fusion. The transform further allows us to map the RGB labels to the IR image frame. Thus, we do not need to label the IR images separately. Not only does this save resources, but labelling the IR images directly can be challenging for the annotators, as the distinction of some classes, such as ice and water, can be unclear. In our dataset we found no reason to label the IR images separately, however, we note that there can be cases where there are clear advantages of doing so. For example, the IR camera is less affected by poor visibility (due to darkness, fog, etc.), or there might be objects which are invisible in the RGB image, yet fully visible in the IR image. Transforming the RGB label to the IR image frame makes little sense in such cases, and will only limit the potential of the IR-based detection. However, our dataset was collected in clear daylight conditions and therefore the given example does not apply.

The sea-ice classification plays an important role in the sea-ice detection. In recent years, deep convolutional neural networks based architectures have replaced the hand-crafting feature extraction techniques with the automated one in image classification. In this project, we consider semantic segmentation for sea-ice classification, which aims to classify each pixel in an image belonging to a predefined class labels.

There exist more than 100 image segmentation methods based on deep neural networks. However, we will focus on few of them for sea-ice classification, particularly, U-Net, a variation of U-Net with Visual Geometry Group (VGG), pyramid scheme parsing network (PSPNet), and Detectron2 with its code-ready implementation. We will investigate these methods and choose the best one suited for sea-ice detection.

Based on the study of different ML architectures, the Detectron2 network by Facebook AI research (FAIR) group was selected to be used. The Detectron2 is a platform for object detection, panoptic segmentation, Densepose, etc. It comprises of a large collection of pre-trained models, however, we only focus on one of its architectures, the `MaskRCNN-R-50-FPN-3x`. This architecture provides the segmentation boundary of object of interest with its confidence value for each segment in image and was deemed the most suitable for ice classifications. We utilize an existing implementation of the architecture with a COCO-panoptic-segmentation (Common Objects in COntext) and pre-trained weights which was then retrained on the Sampo dataset.

In the final step, the semantic segmentation outputs are post-processed to improve the segmentation results. A schematic outline of the post-processing steps is shown in Fig. 1. The post-processing is aimed to reduce false detections, for example ice detections above the horizon, or to filter speckled output frames, for example in cases where there are small miss labelled water patches within icy regions. The post-processing is done on a confidence map containing the confidence of each class: `background`, `ice-track`, `buoy`, `ice`, and `water`. The confidence map is normalized and based on background output S_h, a horizon threshold $\hat{S}_h$ is obtained. A horizon level confidence with image height is obtained by averaging horizontally in $\hat{S}_h$ and a horizon level is obtained with 90% threshold. All detections above horizon level are set to background.

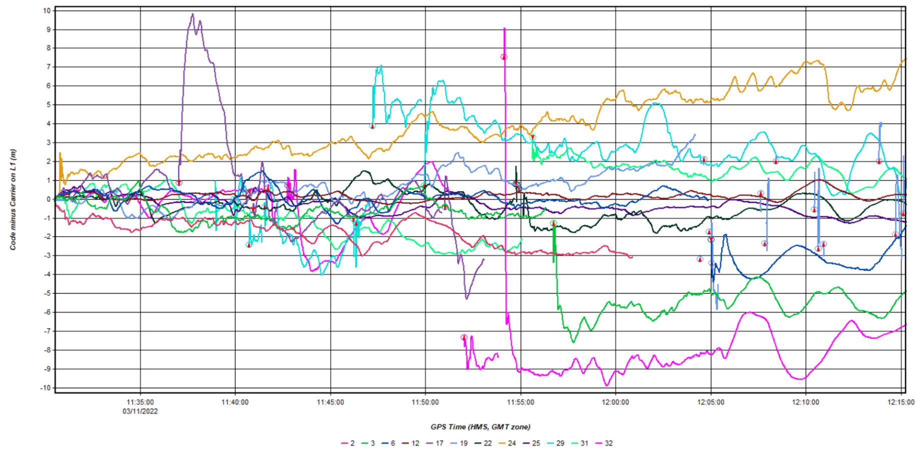

Figure 2: Example of the GNSS observations from the start of the Kemi campaign, showing the CMC measurements aboard the icebreaker Sampo.

From the processed output obtained in previous step, the pixel segmentation map is obtained based on maximum confidence for the class based on a *argmax* function along class axis. It is filtered for blobs of size 1% in classes `ice-track`, `ice`, and `water`, and then holes in these class segmentations are filled using binary hole filling.

2 Results on the GNSS multipath analysis

One of the major sources of error in GNSS measurements is the effect of multipath, caused by signal reflections from nearby obstacles. The reflected signals typically have different power level, code delay, and suffer from changes in carrier phase and frequency. These effects result in distortions between the replica signal generated by the receiver and the true GNSS signal, which can cause tens of meters of error in both pseudorange and code phase estimates.

We use the well defined Code-minus-Carrier (CMC) method, where the code multipath error is estimated by subtracting carrier phase measurements from corresponding pseudoranges. However, because the multipath error on the pseudorange estimate is typically significantly larger compared to the carrier phase error, the CMC method is mostly related to the pseudorange multipath. The subtraction removes the effect of non-dispersive systematic errors such as satellite and receiver clock errors, possible orbital errors, and delays caused by the troposphere.

Due to poor ice coverage in front of Helsinki during the winter 2021 - 2022, the GNSS data collection was done only during the Kemi field campaign on board the icebreaker Sampo. Because the MS Megastar was chosen as our primary test platform, installation of GNSS antenna had to be done on a short notice before the icebreaker Sampo set sail. Thus, the location of the antenna was not optimal, with the mast and the bridge of the ship causing shadowing.

The results of our multipath analysis, shown in Figs. 2 and 3 show that the residuals between the carrier and code phases are on average between 2-3 meters for the first $\sim$15 minutes of the test campaign, after which the difference increases sharply for four satellites (G3, G24, G29, G32). The increase in the CMC is likely caused by the location of the antenna, and the change of the ship trajectory. As the ship slowly changes its heading, the mast and the bridge of the ship are slowly causing varying shadowing on the antenna. The shadowing causes the fast jumps seen in the CMC and the increase in cycle slips indicated with the circles with arrows (arrow pointing down, signal lost, arrow pointing up, signal regained).

The multipath tests were carried out by tilting the antenna towards the ice surface. As expected, because of the tilt angle, GPS satellites G2 and 19 were lost from the positioning solution, but satellites G20 and G32 have been added. For both new satellites, a clear variation in the CMC values is seen, especially for the satellite G20 reaching values as high as $\sim$ 23 meters. However, for the majority of the GPS satellites, the CMC is clearly under 1 meter level for the test period.

Considering the results obtained here, and the sub-optimal placement of the GNSS antenna aboard the icebreaker Sampo, it is likely that GNSS multipath effects from sea-ice are not a major concern. The sea-ice does reflect the GNSS signals, but an antenna placed properly in the mast of the ship and equipped with a choke ring should provide enough mitigation. However, further dedicated research campaign would be needed to define how much multipath from sea-ice affects GNSS measurements.

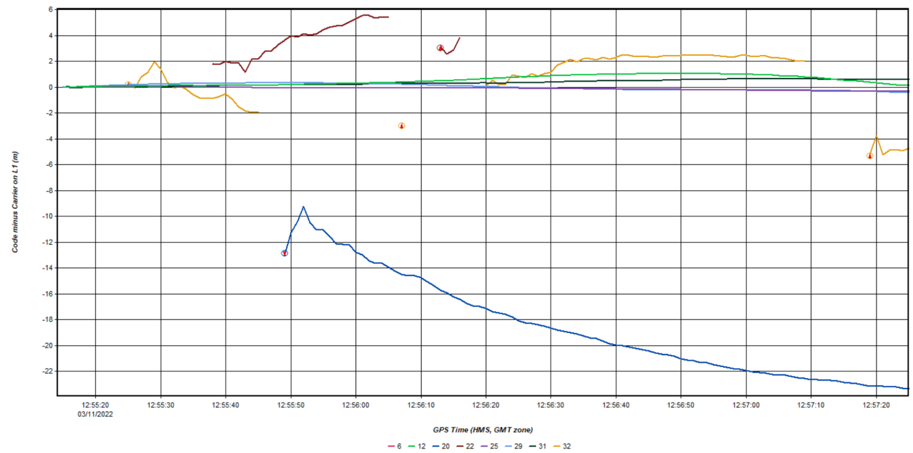

Figure 3: Example of the GNSS multipath observations during the Kemi campaign, showing the CMC measurements with the antenna tilted towards the sea-ice while the Sampo icebreaker was stationary.

3 Results on the fused IR/RGB prediction and ML methods

The benefit of adding IR cameras to vessels is clearly demonstrated in Fig. 4, where the first two panels on the first row show the raw RGB and IR images, respectively. The ice-track is clearly visible in both images, but the RGB frame shows more details, whereas the IR is smoother. In this example, the IR image is better for ice-track detection, but the RGB image provides more information on the type of ice the vessels is moving. Both IR and RGB image also have two buoys visible right above the ice-track, but the annotated labels, third panel on the top row has only ice and ice-track labelled. Furthermore, the ice-track label is missing a significant part of the actual ice-track. Our data set already has a low number of images with buoys, and combined with the missing labels in some images, has resulted in the low accuracy in the network detection accuracy. The problems with labelling are also affecting the general performance of the network.

The detector outputs, middle row in Fig. 4, show that the RGB output has noisy and broken detection of ice tracks in regions that should be ice. The IR output has less noisy detections and combining the RGB in IR data has reduced this further. However, the ice-track detection is better in the RGB output, which is related to the lower number of available IR data. The post-processing has greatly reduced the noisy detections in the RGB output, but in the IR output the post-processing has replaced the ice-track with background label. The post-processing steps were not optimized and the used methods were simplified, but we expect that with more detailed optimization, and with more IR data, the post processing will improve considerably the outputs from the network.

While our results serve as a promising proof of concept, the dataset size and quality are likely the main bottleneck for greater performance. A large scale data pipeline similar to the ones used by car manufacturers would be needed. Car manufacturers interested in autonomous driving have several sensors (cameras etc.) installed in each of their vehicles. The sensors collect data continuously, and transmit it to an expert in-house data annotation team, from which the data is moved to model training and iteration. However, such a data pipeline is not possible to achieve in a research project of this scale.

The usefulness of the results for winter navigation can be viewed from two different perspectives. Firstly from the perspective of an autonomous navigation system that could utilise the resulting output data for autonomous decision making, and secondly from the perspective of a human operator who could use the resulting output data as assisting in traditional navigation. In both cases we gauge that the data this type of a system generates, already at this trial stage, can be useful for safe navigation in winter and ice conditions.

For a human operator, on a conventional vessel, the shown results can assist in forming a better awareness of the surrounding ice situation. The results provided by the IR camera can be seen as especially useful on cargo vessels with the bridge at the aft, since navigation in ice at night requires powerful floodlights that have the disadvantage of highlighting possible snow fall and deck constructions and thus heavily reducing the visibility. By using IR cameras installed in the fore mast or fore part of the vessel, the operator can benefit from a better view of the ice field in front of the vessel without unnecessarily highlighting the surroundings with floodlights. For a human operator the various issues in the output data, i.e. small speckles in ice fields or false positive detections above the horizon or in archipelago waters, can easily be filtered out by the operator and only the relevant ice track data is used for decision making and for steering the vessel.

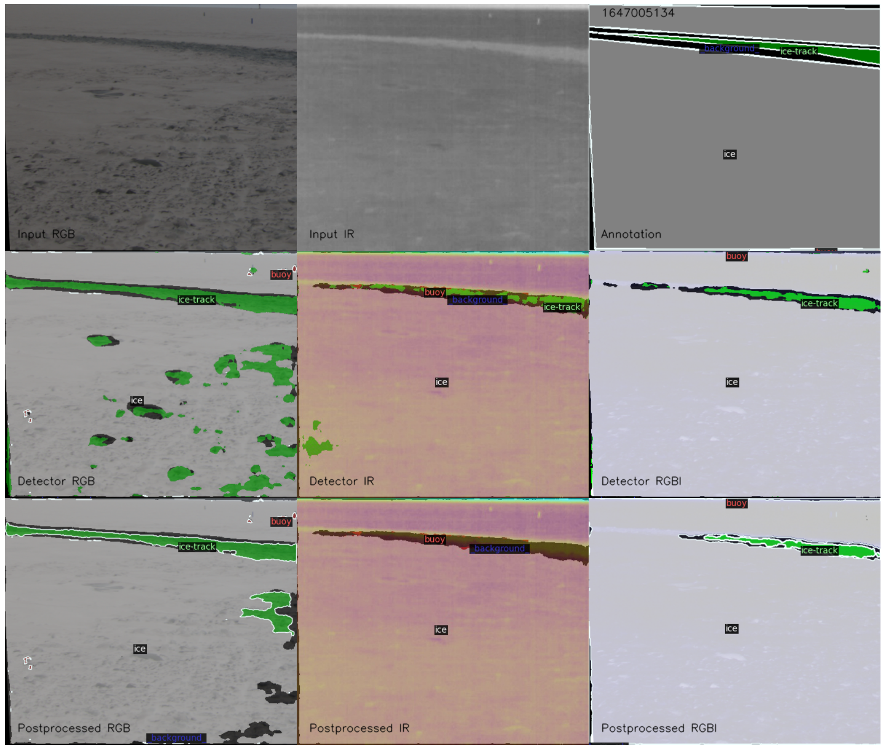

Figure 4: Comparison between the raw and fused camera data and the machine learning output. The first row shows the input RGB and IR images, and the machine learning labels. The second row shows the detector output for the RGB, IR, and fused RGBI images, and the third row shows the post-processed RGB, IR, and fused RGBI images.

4 Conclusions

The ENHANCE project has shown that the fusion of visible and infrared camera information can be used to increase situational awareness of the surrounding ice conditions and that ML methods can be used to detect sea ice and especially ice-tracks with good accuracy. However, further testing and development is needed before the system developed during the ENHANCE project can be deployed to a fully autonomous vessel. The diversity of the available data is essential for good AI performance, otherwise it may fail to generalize image spatio-chromatic information according to desired labels. However, the number of publicly available data for sea ice studies is limited.

Our results indicate that Infrared cameras can increase the operational stability of any vessel, manned or autonomous, during night-time or low-light conditions, but this was not tested during the project. However, it was shown that pure IR solution can detect ice tracks with as high accuracy as RGB only solution, thus underlining the benefits of IR cameras. In addition to the camera fusion, a GNSS multipath test campaign was performed. Based on our measurements, we do not expect that multipath is a major issue for autonomous navigation in icy conditions. However, a specialised test campaign would be required for a detailed study of multipath effects caused by sea ice to vessels navigating in icy conditions.

The results presented in this project can be improved by increasing the available data, from both RGB and IR cameras, as well as available GNSS measurements. For future projects, focusing on methodology and obtaining detailed scene descriptions for the rapidly growing AI field is advised. Manual labour required to label the camera data for ML training is also challenging and can be extremely time consuming, which requires careful arrangements and backup planning.

The ML approach studied in this project is limited by the data used for training the model, since the model might not be general enough and the system will only work on limited data, or in settings which are representative of that data. For future projects one would need a closed loop evaluation based on how well the ship can navigate based on these detections. However, the method developed during the project has been a successfully proof of a concept of using AI and ML based approach to increase the situational averseness while navigating in icy conditions, but more research is needed before it can be used for truly autonomous navigation.